



Intel® Cluster Checker Reference Manual

Version 3.1.2

January 15, 2016

Disclaimer and Legal Information

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Requires a system with a 64-bit enabled processor, chipset, BIOS and software. Performance will vary depending on the specific hardware and software you use. Consult your PC manufacturer for more information. For more information, visit <http://www.intel.com/info/em64t>

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>

Intel, the Intel logo, the Intel Inside logo, Xeon, and Xeon Phi are trademarks of Intel Corporation in the U.S. and/or other countries.

Optimization Notice

Intel compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

* Other names and brands may be claimed as the property of others.

© 2016 Intel Corporation. All rights reserved.

Contents

1 Analyzer	6
2 Data Providers	14
2.1 hpl_cluster	15
2.2 hpl_pairwise	19
2.3 ibstat	23
2.4 ibv_devinfo	24
2.5 iofzone	24
2.6 imb_pingpong	26
2.7 ofedinfo	30
2.8 opa	31
3 Data Collection	32
A List of Checks	35

Chapter 1

Analyzer

The behavior of Intel® Cluster Checker can be configured through an XML file that is specified as a command line configuration option. The default location of the file is `/opt/intel/clck/3.1.2/etc/clck.xml`

The `config` XML tag can be used to configure the behavior of the knowledge base.

`clck-checks`

Specify the checks to be performed.

XML syntax

```
<config>
  <clck-checks>
    <entry>CHECK</entry>
    <entry>CHECK</entry>
  </clck-checks>
</config>
```

Description

The `clck-checks` tag specifies the checks that should be performed. Only checks that are specified will be performed.

Valid values for `CHECK` are given in Appendix A.

The `entry` tag can be repeated multiple times.

`clck-mode`

Specify the mode under which Intel® Cluster Checker should run.

XML syntax

```
<config>
  <clck-mode>MODE</clck-mode>
</config>
```

Description

This specifies the mode that Intel® Cluster Checker should run under. Please see the User's Guide for more information.

Valid values for MODE are certification, compliance and health.

data-age-threshold

Specify the maximum age of data points, in seconds, before a data point is considered "too old" for relevant analysis.

XML syntax

```
<config>  
  <data-age-threshold>NUMBER</data-age-threshold>  
</config>
```

Description

The data-age-threshold tag specifies the maximum age of data points, in seconds, before a data point is considered "too old" for relevant analysis.

The value should be an integer value greater than 0. The default value is 604800 seconds (1 week).

dgemm-number-of-mad

Specify the number of median absolute deviations (MADs) allowed before a dgemm value is considered an outlier.

XML syntax

```
<config>  
  <dgemm-number-of-mad>NUMBER</dgemm-number-of-mad>  
</config>
```

Description

The dgemm-number-of-mad tag specifies the number of allowed median absolute deviations before a dgemm value is considered an outlier.

The value should be an integer value greater than 0.

dgemm-peak-fraction

Specify the minimum value of the ratio between the measured dgemm performance and theoretical peak value.

XML syntax

```
<config>  
  <dgemm-peak-fraction>NUMBER</dgemm-peak-fraction>  
</config>
```

Description

The `dgemm-peak-fraction` tag specifies the minimum value for the ratio between the measured dgemm performance and the theoretical peak performance. Any value below this will generate a sign.

The value should be an floating point value between 0 and 1.

environment-blacklist

Specify the environment variable patterns that will be ignored for uniformity comparison.

XML syntax

```
<config>
  <environment-blacklist>
    <entry>PATTERN</entry>
    <entry>PATTERN</entry>
  </environment-blacklist>
</config>
```

Description

The `environment-blacklist` tag excludes environment variables from being compared for uniformity across the cluster.

The value within each entry tag is interpreted as a POSIX* matching regular expression. If this value is not a valid POSIX* regular expression, then no filtering will be done.

The entry tag can be repeated multiple times.

Note that to exactly match meta characters, (`^[.*{\\(||?<>]`), they should be escaped.

hpl-number-of-mad

Specify the number of median absolute deviations (MADs) allowed before an HPL value is considered an outlier.

XML syntax

```
<config>
  <hpl-number-of-mad>NUMBER</hpl-number-of-mad>
</config>
```

Description

The `hpl-number-of-mad` tag specifies the number of allowed median absolute deviations before an HPL value is considered an outlier.

The value should be an integer value greater than 0.

icr-version

Specify the Intel® Cluster Ready specification version.

XML syntax

```
<config>  
  <icr-version>VERSION</icr-version>  
</config>
```

Description

The icr-version tag specifies the Intel® Cluster Ready specification version.

The acceptable values are 1.3.1 and 1.4. The default value is 1.4.

imb-pingpong-number-of-mad

Specify the number of median absolute deviations (MADs) allowed before an HPL value is considered an outlier.

XML syntax

```
<config>  
  <imb-pingpong-number-of-mad>NUMBER</imb-pingpong-number-of-mad>  
</config>
```

Description

The imb-pingpong-number-of-mad tag specifies the number of allowed median absolute deviations before a PingPong latency or bandwidth value is considered an outlier.

The value should be an integer value greater than 0.

iozone-number-of-mad

Specify the number of median absolute deviations (MADs) allowed before an iozone value is considered an outlier.

XML syntax

```
<config>  
  <iozone-number-of-mad>NUMBER</iozone-number-of-mad>  
</config>
```

Description

The iozone-number-of-mad tag specifies the number of allowed median absolute deviations before an iozone value is considered an outlier.

The value should be an integer value greater than 0.

kernel-blacklist

Specify the kernel parameter patterns that will be ignored for uniformity comparisons.

XML syntax

```
<config>
  <kernel-blacklist>
    <entry>PATTERN</entry>
    <entry>PATTERN</entry>
  </kernel-blacklist>
</config>
```

Description

The kernel-blacklist tag excludes kernel parameters from being compared for uniformity across the cluster.

The value within each entry tag is interpreted as a POSIX* matching regular expression. If this value is not a valid POSIX* regular expression, then no filtering will be done.

The entry tag can be repeated multiple times.

Note that to exactly match meta characters, (`^[.*${\()|?<>}`), they should be escaped.

lshw-blacklist

Specify the lshw output patterns that will be ignored for uniformity comparison.

XML syntax

```
<config>
  <lshw-blacklist>
    <entry>PATTERN</entry>
    <entry>PATTERN</entry>
  </lshw-blacklist>
</config>
```

Description

The lshw-blacklist tag excludes lshw items from being compared for uniformity across the cluster.

The value within each entry tag is interpreted as a POSIX* matching regular expression. If this value is not a valid POSIX* regular expression, then no filtering will be done.

The entry tag can be repeated multiple times.

Note that to exactly match meta characters, (`^[.*${\()|?<>}`), they should be escaped.

lshw-uniformity-threshold

Specify the threshold ratio for checking the uniformity of lshw entries.

XML syntax

```
<config>
  <lshw-uniformity-threshold>NUMBER</lshw-uniformity-threshold>
</config>
```

Description

The lshw uniformity check determines whether each lshw entry is uniform across the cluster. If the percentage of nodes that share the same lshw entry value is above the value specified for the lshw-uniformity-threshold tag, then that value is considered uniform in that cluster. If the percentage of nodes that share the same lshw entry value is below the uniformity threshold, then a sign is generated.

The value should be an floating point value between 0 and 1.

memory-uniformity-threshold

Specify the maximum allowable deviation, in bytes, from the median memory size before a memory size is considered non-uniform.

XML syntax

```
<config>
  <memory-uniformity-threshold>NUMBER</memory-uniformity-
    threshold>
</config>
```

Description

The memory-uniformity-threshold tag specifies the maximum allowable deviation, in bytes, from the median memory size before a memory size is considered non-uniform.

Any value greater than 0 can be used for this tag. The default value is 268435456 bytes (256 MB).

ntp-offset-threshold

Specify the maximum offset value that an NTP peer is allowed to have.

XML syntax

```
<config>
  <ntp-offset-threshold>NUMBER</ntp-offset-threshold>
</config>
```

Description

The ntp-offset-threshold tag specifies the maximum offset value that an NTP peer is allowed to have before a sign is generated.

Any floating point value can be used for this tag.

outlier-max-median-mad-dist

Specify the maximum distance between the median and median absolute deviation (MAD) for the MAD outlier algorithm to be used.

XML syntax

```
<config>
  <outlier-max-median-mad-dist>NUMBER</outlier-max-median-mad-
    dist>
</config>
```

Description

The outlier-max-median-mad-dist tag specifies the maximum allowable distance, in orders of magnitude, between the median and the median absolute deviation. If the allowable distance is exceeded, then the MAD outlier algorithm is disabled and a fallback algorithm (controlled by the outlier-median-pct tag) is used for outlier checks.

The following describes the test controlled by the outlier-max-median-mad-dist tag:

```
if ( |median - MAD| < 10^outlier-max-median-mad-dist )
  then <use MAD outlier algorithm>
  else <use fallback outlier algorithm>
```

Any value greater than 0 can be used for this tag. The default value is 2.5.

outlier-median-pct

Percentage of the median used to calculate outliers by the fallback algorithm.

XML syntax

```
<config>
  <outlier-median-pct>NUMBER</outlier-median-pct>
</config>
```

Description

The outlier-median-pct tag determines the distance from the median that a sample value is allowed to be before it is considered an outlier in the fallback outlier algorithm. The outlier-median-pct value is divided by 100 and multiplied by the median to get an allowable distance. If the sample value is further away from the median than the allowable distance, the sample value is considered an outlier.

The following describes the fallback outlier algorithm controlled by the outlier-median-pct tag:

```
if ( |median - sample_value| >
  (median * (outlier-median-pct / 100) )
  then <the sample_value is an outlier>
  else <the sample_value is not an outlier>
```

Any value between 0 and 100 can be used. The default value is 5.

rpm-uniformity-threshold

Specify the threshold ratio for checking the uniformity of rpm entries.

XML syntax

```
<config>  
  <rpm-uniformity-threshold>NUMBER</rpm-uniformity-threshold>  
</config>
```

Description

The rpm uniformity check determines whether each rpm file installed on a node is uniform across the cluster. If the percentage of nodes that share the same rpm file is above the value specified for the rpm-uniformity-threshold tag, then that rpm file is considered uniform on the cluster. If the percentage of nodes that share the same rpm file is below the uniformity threshold, then a sign is generated.

The value should be a floating point value between 0 and 1.

storage-max-used-pct

The maximum percentage of space that can be used on a partition.

XML syntax

```
<config>  
  <storage-max-used-pct>NUMBER</storage-max-used-pct>  
</config>
```

Description

The storage-max-used-pct tag specifies the maximum percentage of space that can be used on a disk partition. If the percentage is exceeded on a disk partition then a sign is emitted.

Any value between 0 and 100 can be used for this tag. The default value is 85.

stream-number-of-mad

Specify the number of median absolute deviations (MADs) allowed before a stream value is considered an outlier.

XML syntax

```
<config>  
  <stream-number-of-mad>NUMBER</stream-number-of-mad>  
</config>
```

Description

The stream-number-of-mad tag specifies the number of allowed median absolute deviations before a stream value is considered an outlier.

The value should be an integer value greater than 0.

Chapter 2

Data Providers

The data provider configuration mechanism is environment variables. The name space `CLCK_PROVIDER_PROVIDER-NAME` is reserved for each provider, where the `PROVIDER-NAME` is the name of the corresponding provider. Each provider can define its own variables within its name space, and each provider is responsible for reading its own environment variables.

While these variables can be set in the users shell initialization files, e.g., `.bashrc`, the recommended way is to define these variables in the configuration file that is read and propagated by `clck-collect` and `clckd`.

The XML configuration file hides the raw environment variable names. Instead, the environment variable gets constructed by combining the XML tag names.

For example, the following construct -

```
<configuration>
  <provider>
    <hpl_cluster>
      <fabrics>value</fabrics>
    </hpl_cluster>
  </provider>
</configuration>
```

will generate the environment variable `CLCK_PROVIDER_HPL_CLUSTER_FABRICS`, which would then be made available to the `HPL_CLUSTER` provider.

Note, the value is enclosed within double quotes when passed to the provider. Consequently, if the value provided in the XML contains double quotes, it will be ignored and the environment variable will not be constructed.

Additionally, XML configuration files can also be used to enable running the same provider multiple times with disjoint sets of configuration options. For example, the following XML configuration would run the `HPL_CLUSTER` twice, once with `CLCK_PROVIDER_HPL_CLUSTER_FABRICS=shm:tcp` and once with `CLCK_PROVIDER_HPL_CLUSTER_FABRICS=shm:dapl`

```
<configuration>
```

```

<provider>
  <hpl_cluster>
    <optionset id="1">
      <fabrics>shm:tcp</fabrics>
    </optionset>

    <optionset id="2">
      <fabrics>shm:dapl</fabrics>
    </optionset>
  </hpl_cluster>
</provider>
</configuration>

```

The “id” attribute is required. For each run, the value of the corresponding “id” attribute will be stored in the database under the “OptionID” column. The value of this attribute is purely arbitrary and is not meaningful other than to distinguish one set of options from another in the output columns of the database. However, the value must not contain tabs, newlines, parentheses, ampersands, vertical bars, the symbol ‘<’, the symbol ‘~’, spaces, semicolons or double quotes.

If no “optionset” tags are specified, the “OptionID” column will contain the string “default”.

Additionally, if any tags are specified outside of “optionset” tags, the corresponding environment variables will be made available for every single “optionset” specified.

2.1 hpl_cluster

CLK_PROVIDER_HPL_CLUSTER_FABRICS

Select the network fabric to be used.

Environment variable syntax

CLK_PROVIDER_HPL_CLUSTER_FABRICS=value

XML syntax

```

<configuration>
  <provider>
    <hpl_cluster>
      <fabrics>value</fabrics>
    </hpl_cluster>
  </provider>
</configuration>

```

Arguments

value

The network fabric to use.

Description

This value directly maps to the `I_MPI_FABRICS` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate fabric.

CLCK_PROVIDER_HPL_CLUSTER_DAPL_PROVIDER

Select the DAPL provider to use.

Environment variable syntax

`CLCK_PROVIDER_HPL_CLUSTER_DAPL_PROVIDER=value`

XML syntax

```
<configuration>
  <provider>
    <hpl_cluster>
      <dapl_provider>value</dapl_provider>
    </hpl_cluster>
  </provider>
</configuration>
```

Arguments

value

The DAPL provider to use.

Description

This value directly maps to the `I_MPI_DAPL_PROVIDER` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate DAPL provider.

CLCK_PROVIDER_HPL_CLUSTER_DAPL_PROVIDER_LIST

Select the DAPL providers to load.

Environment variable syntax

`CLCK_PROVIDER_HPL_CLUSTER_DAPL_PROVIDER_LIST=value`

XML syntax

```
<configuration>
  <provider>
    <hpl\_cluster>
      <dapl_provider_list>value</dapl_provider_list>
    </hpl\_cluster>
  </provider>
</configuration>
```



```

    </hpl\_cluster>
  </provider>
</configuration>

```

Arguments

value

The DAPL providers to load.

Description

This value directly maps to the `I_MPI_DAPL_PROVIDER_LIST` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate DAPL providers.

CLCK_PROVIDER_HPL_CLUSTER_OFA_ADAPTER_NAME

Select the OFA adapter name to use.

Environment variable syntax

CLCK_PROVIDER_HPL_CLUSTER_OFA_ADAPTER_NAME=value

XML syntax

```

<configuration>
  <provider>
    <hpl_cluster>
      <ofa_adapter_name>value</ofa_adapter_name>
    </hpl_cluster>
  </provider>
</configuration>

```

Arguments

value

The OFA adapter name to use.

Description

This value directly maps to the `I_MPI_OFA_ADAPTER_NAME` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate OFA adapter name.

CLCK_PROVIDER_HPL_CLUSTER_TCP_NETMASK

Select the TCP netmask value to use.

Environment variable syntax

CLK_PROVIDER_HPL_CLUSTER_TCP_NETMASK=value

XML syntax

```
<configuration>
  <provider>
    <hpl_cluster>
      <tcp_netmask>value</tcp_netmask>
    </hpl_cluster>
  </provider>
</configuration>
```

Arguments

value

The TCP netmask value to use.

Description

This value directly maps to the I_MPI_TCP_NETMASK Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate OFA adapter name.

CLK_PROVIDER_HPL_CLUSTER_TMI_PROVIDER

Select the TMI provider to use.

Environment variable syntax

CLK_PROVIDER_HPL_CLUSTER_TMI_PROVIDER=value

XML syntax

```
<configuration>
  <provider>
    <hpl_cluster>
      <tmi_provider>value</tmi_provider>
    </hpl_cluster>
  </provider>
</configuration>
```

Arguments

value

The TMI provider to use.

Description

This value directly maps to the `I_MPI_TMI_PROVIDER` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate TMI provider.

CLCK_PROVIDER_HPL_CLUSTER_OPTIONS

Any additional desired options to be provided to the `mpirun` command.

Environment variable syntax

`CLCK_PROVIDER_HPL_CLUSTER_OPTIONS=value`

XML syntax

```
<configuration>
  <provider>
    <hpl_cluster>
      <options>value</options>
    </hpl_cluster>
  </provider>
</configuration>
```

Arguments

value

Any options that are passed verbatim to the `mpirun` command.

Description

The value of this tag is passed as is to the `mpirun` command.

2.2 hpl_pairwise

CLCK_PROVIDER_HPL_PAIRWISE_FABRICS

Select the network fabric to be used.

Environment variable syntax

`CLCK_PROVIDER_HPL_PAIRWISE_FABRICS=value`

XML syntax

```
<configuration>
  <provider>
    <hpl_pairwise>
      <fabrics>value</fabrics>
    </hpl_pairwise>
  </provider>
</configuration>
```

Arguments

value

The network fabric to use.

Description

This value directly maps to the I_MPI_FABRICS Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate fabric.

CLICK_PROVIDER_HPL_PAIRWISE_DAPL_PROVIDER

Select the DAPL provider to use.

Environment variable syntax

CLICK_PROVIDER_HPL_PAIRWISE_DAPL_PROVIDER=value

XML syntax

```
<configuration>
  <provider>
    <hpl_pairwise>
      <dapl_provider>value</dapl_provider>
    </hpl_pairwise>
  </provider>
</configuration>
```

Arguments

value

The DAPL provider to use.

Description

This value directly maps to the I_MPI_DAPL_PROVIDER Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate DAPL provider.

CLICK_PROVIDER_HPL_PAIRWISE_DAPL_PROVIDER_LIST

Select the DAPL providers to load.

Environment variable syntax

CLICK_PROVIDER_HPL_PAIRWISE_DAPL_PROVIDER_LIST=value

XML syntax

```
<configuration>
  <provider>
    <hpl_pairwise>
      <dapl_provider_list>value</dapl_provider_list>
    </hpl_pairwise>
  </provider>
</configuration>
```

Arguments

value

The DAPL providers to load.

Description

This value directly maps to the `I_MPI_DAPL_PROVIDER_LIST` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate DAPL providers.

CLCK_PROVIDER_HPL_PAIRWISE_OFA_ADAPTER_NAME

Select the OFA adapter name to use.

Environment variable syntax

CLCK_PROVIDER_HPL_PAIRWISE_OFA_ADAPTER_NAME=value

XML syntax

```
<configuration>
  <provider>
    <hpl_pairwise>
      <ofa_adapter_name>value</ofa_adapter_name>
    </hpl_pairwise>
  </provider>
</configuration>
```

Arguments

value

The OFA adapter name to use.

Description

This value directly maps to the `I_MPI_OFA_ADAPTER_NAME` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate OFA adapter name.

CLCK_PROVIDER_HPL_PAIRWISE_TCP_NETMASK

Select the TCP netmask value to use.

Environment variable syntax

CLCK_PROVIDER_HPL_PAIRWISE_TCP_NETMASK=value

XML syntax

```
<configuration>
  <provider>
    <hpl_pairwise>
      <tcp_netmask>value</tcp_netmask>
    </hpl_pairwise>
  </provider>
</configuration>
```

Arguments

value

The TCP netmask value to use.

Description

This value directly maps to the I_MPI_TCP_NETMASK Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate OFA adapter name.

CLCK_PROVIDER_HPL_PAIRWISE_TMI_PROVIDER

Select the TMI provider to use.

Environment variable syntax

CLCK_PROVIDER_HPL_PAIRWISE_TMI_PROVIDER=value

XML syntax

```
<configuration>
  <provider>
    <hpl_pairwise>
      <tmi_provider>value</tmi_provider>
    </hpl_pairwise>
  </provider>
</configuration>
```

Arguments

value

The TMI provider to use.

Description

This value directly maps to the `I_MPI_TMI_PROVIDER` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate TMI provider.

CLCK_PROVIDER_HPL_PAIRWISE_OPTIONS

Any additional desired options to be provided to the `mpirun` command.

Environment variable syntax

`CLCK_PROVIDER_HPL_PAIRWISE_OPTIONS=value`

XML syntax

```
<configuration>
  <provider>
    <hpl_pairwise>
      <options>value</options>
    </hpl_pairwise>
  </provider>
</configuration>
```

Arguments

value

Any options that are passed verbatim to the `mpirun` command.

Description

The value of this tag is passed as is to the `mpirun` command.

2.3 ibstat

CLCK_PROVIDER_IBSTAT_BINARY

Select the location of the `ibstat` binary.

Environment variable syntax

`CLCK_PROVIDER_IBSTAT_BINARY=value`

XML syntax

```
<configuration>
  <provider>
    <ibstat>
      <binary>value</binary>
    </ibstat>
  </provider>
</configuration>
```

Arguments

value

The path to the ibstat binary.

Description

This value sets the path to the ibstat binary.

When this variable is not set, /usr/sbin/ibstat will be used.

2.4 ibv_devinfo

CLK_PROVIDER_IBV_DEVINFO_BINARY

Select the location of the ibv_devinfo binary.

Environment variable syntax

CLK_PROVIDER_IBV_DEVINFO_BINARY=value

XML syntax

```
<configuration>
  <provider>
    <ibv_devinfo>
      <binary>value</binary>
    </ibv_devinfo>
  </provider>
</configuration>
```

Arguments

value

The path to the ibv_devinfo binary.

Description

This value sets the path to the ibv_devinfo binary.

When this variable is not set, /usr/bin/ibv_devinfo will be used.

2.5 iozone

CLK_PROVIDER_IOZONE_FILESIZE

Select the size of the file used to test.

Environment variable syntax

CLK_PROVIDER_IOZONE_FILESIZE=value

XML syntax

```
<configuration>
  <provider>
    <iozone>
      <filesize>value</filesize>
    </iozone>
  </provider>
</configuration>
```

Arguments

value

The file size to use, in Kbytes.

Description

This value sets the file size of the temporary file used by the benchmark.

When this variable is not set, 65536 will be used for the file size.

CLK_PROVIDER_IOZONE_RECSIZE

Select the record size used to test.

Environment variable syntax

CLK_PROVIDER_IOZONE_RECSIZE=value

XML syntax

```
<configuration>
  <provider>
    <iozone>
      <recsize>value</recsize>
    </iozone>
  </provider>
</configuration>
```

Arguments

value

The record size to use, in Kbytes.

Description

This value sets the record size used by the benchmark.

When this variable is not set, 16384 will be used for the record size.

CLK_PROVIDER_IOZONE_WORKDIR

Select the directory where a temporary file will be created.

Environment variable syntax

CLK_PROVIDER_IOZONE_WORKDIR=value

XML syntax

```
<configuration>
  <provider>
    <iozone>
      <workdir>value</workdir>
    </iozone>
  </provider>
</configuration>
```

Arguments

value

The location of the directory where a temporary file will be created by the benchmark.

Description

This value sets the directory where a temporary file will be created by the benchmark. It is recommended to set this value to local file system.

When this variable is not set, /tmp will be used for the temporary file location.

2.6 imb_pingpong

CLK_PROVIDER_IMB_PINGPONG_FABRICS

Select the network fabric to be used.

Environment variable syntax

CLK_PROVIDER_IMB_PINGPONG_FABRICS=value

XML syntax

```
<configuration>
  <provider>
    <imb_pingpong>
      <fabrics>value</fabrics>
    </imb_pingpong>
  </provider>
</configuration>
```

Arguments

value

The network fabric to use.

Description

This value directly maps to the I_MPI_FABRICS Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate fabric.

CLCK_PROVIDER_IMB_PINGPONG_DAPL_PROVIDER

Select the DAPL provider to use.

Environment variable syntax

CLCK_PROVIDER_IMB_PINGPONG_DAPL_PROVIDER=value

XML syntax

```
<configuration>
  <provider>
    <imb_pingpong>
      <dapl_provider>value</dapl_provider>
    </imb_pingpong>
  </provider>
</configuration>
```

Arguments

value

The DAPL provider to use.

Description

This value directly maps to the I_MPI_DAPL_PROVIDER Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate DAPL provider.

CLCK_PROVIDER_IMB_PINGPONG_DAPL_PROVIDER_LIST

Select the DAPL providers to load.

Environment variable syntax

CLCK_PROVIDER_IMB_PINGPONG_DAPL_PROVIDER_LIST=value

XML syntax

```
<configuration>
  <provider>
    <imb\_pingpong>
      <dapl_provider_list>value</dapl_provider_list>
    </imb\_pingpong>
  </provider>
</configuration>
```

```

    </imb\_pingpong>
  </provider>
</configuration>

```

Arguments

value

The DAPL providers to load.

Description

This value directly maps to the I_MPI_DAPL_PROVIDER_LIST Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate DAPL providers.

CLCK_PROVIDER_IMB_PINGPONG_OFA_ADAPTER_NAME

Select the OFA adapter name to use.

Environment variable syntax

CLCK_PROVIDER_IMB_PINGPONG_OFA_ADAPTER_NAME=value

XML syntax

```

<configuration>
  <provider>
    <imb_pingpong>
      <ofa_adapter_name>value</ofa_adapter_name>
    </imb_pingpong>
  </provider>
</configuration>

```

Arguments

value

The OFA adapter name to use.

Description

This value directly maps to the I_MPI_OFA_ADAPTER_NAME Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate OFA adapter name.

CLCK_PROVIDER_IMB_PINGPONG_TCP_NETMASK

Select the TCP netmask value to use.

Environment variable syntax

CLCK_PROVIDER_IMB_PINGPONG_TCP_NETMASK=value

XML syntax

```
<configuration>
  <provider>
    <imb_pingpong>
      <tcp_netmask>value</tcp_netmask>
    </imb_pingpong>
  </provider>
</configuration>
```

Arguments

value

The TCP netmask value to use.

Description

This value directly maps to the I_MPI_TCP_NETMASK Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate OFA adapter name.

CLCK_PROVIDER_IMB_PINGPONG_TMI_PROVIDER

Select the TMI provider to use.

Environment variable syntax

CLCK_PROVIDER_IMB_PINGPONG_TMI_PROVIDER=value

XML syntax

```
<configuration>
  <provider>
    <imb_pingpong>
      <tmi_provider>value</tmi_provider>
    </imb_pingpong>
  </provider>
</configuration>
```

Arguments

value

The TMI provider to use.

Description

This value directly maps to the `I_MPI_TMI_PROVIDER` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate TMI provider.

CLK_PROVIDER_IMB_PINGPONG_OPTIONS

Any additional desired options to be provided to the `mpirun` command.

Environment variable syntax

`CLK_PROVIDER_IMB_PINGPONG_OPTIONS=value`

XML syntax

```
<configuration>
  <provider>
    <imb_pingpong>
      <options>value</options>
    </imb_pingpong>
  </provider>
</configuration>
```

Arguments

value

Any options that are passed verbatim to the `mpirun` command.

Description

The value of this tag is passed as is to the `mpirun` command.

2.7 ofedinfo

CLK_PROVIDER_OFEDINFO_BINARY

Select the location of the `ofedinfo` binary.

Environment variable syntax

`CLK_PROVIDER_OFEDINFO_BINARY=value`

XML syntax

```
<configuration>
  <provider>
    <ofedinfo>
      <binary>value</binary>
    </ofedinfo>
  </provider>
</configuration>
```

Arguments

value

The path to the ofedinfo binary.

Description

This value sets the path to the ofedinfo binary. When this variable is not set, /usr/bin/ofed_info will be used.

2.8 opa

CLK_PROVIDER_OPA_PATH

Select the path to the installed Intel® OPA diagnostic tools.

Environment variable syntax

CLK_PROVIDER_OPA_PATH=value

XML syntax

```
<configuration>
  <provider>
    <opa>
      <path>value</path>
    </opa>
  </provider>
</configuration>
```

Arguments

value

The path to the Intel® OPA diagnostic tools.

Description

This value sets the path to the Intel® OPA diagnostic tools.

When this variable is not set, /usr/sbin/ will be used.

Chapter 3

Data Collection

The following options alter the behavior of data collection.

CLCK_COLLECT_DATABASE_BUSY_TIMEOUT

Specify the amount of time to wait for a database lock to become available.

Environment variable syntax

`CLCK_COLLECT_DATABASE_BUSY_TIMEOUT=value`

Arguments

value

The number of milliseconds to wait for a database lock to become available before giving up. The value must be greater than 0. The default value is 60,000 milliseconds.

Description

When inserting a new row into the database, the database is locked and any concurrent write attempts are prevented. This value specifies the amount of time that the concurrent write(s) should wait for the database to be unlocked before giving up. If the timeout expires and the database is still locked, the concurrent write(s) will not be successful and the data will be lost.

CLCK_COLLECT_DATABASE_CLOSE_DELAY

Specify the amount of time to wait after data collection has finished for data to arrive.

Environment variable syntax

`CLCK_COLLECT_DATABASE_CLOSE_DELAY=value`

Arguments

value

The number of seconds to wait after data collection has finished for any remaining data to be accumulated. The value must be greater than 0. The default value is 1 second.

Description

All data that is in the accumulate queue will always be written to the database, but some data may still be “on the wire” when data collection has finished. This option provides a method to wait an additional amount of time for data to be received by the accumulate server before exiting. Clusters with very slow networks or a very large number of nodes may need to increase this value from the default.

This option is only recognized when using on-demand data collection.

CLKK_COLLECT_DATABASE_VFS_MODULE

Specify the SQLite* VFS module.

Environment variable syntax

CLKK_COLLECT_DATABASE_VFS_MODULE=value

Arguments

value = unix

Uses POSIX advisory locks when locking the database. Note that the implementation of POSIX advisory locks on some filesystems, e.g., NFS, is incomplete and/or buggy. This value should usually only be selected when the database is located on a local filesystem.

value = unix-dotfile

Uses dot-file locking when locking the database. This value usually works around filesystem implementation issues related to POSIX advisory locks. This is the default value.

value = unix-excl

Obtains and holds an exclusive lock on the database file. All concurrent database operations will be prevented while the lock is held. This value may help in the event of database errors during data collection or if collected data is missing from the database.

value = unix-none

No locking is used. This option should only be used if there is a guarantee that only a single writer will modify the database at any given time. Otherwise, this value can easily result in database corruption if two or more processes are writing to the database concurrently.

Description

The SQLite* OS interface layer, or “VFS”, can be selected at runtime. The VFSes differ primarily in the way they handle file locking. See <http://www.sqlite.org/vfs.html> for more information.

This option is only recognized when using on-demand data collection.

Appendix A

List of Checks

Check	Description	Default State
all_to_all	IP address consistency	Enabled
cpu	Intel® Cluster Ready CPU compliance	Enabled
datconf	InfiniBand* DAPL configuration	Enabled
dgemm	Floating point performance	Enabled
environment	Environment variables	Enabled
ethernet	Ethernet driver uniformity and wellness	Enabled
heartbeat	Verify data is recent	Enabled
hpl	High Performance Linpack*	Enabled
icr_cluster	Intel® Cluster Ready minimum node count compliance	Enabled
icr_version	Intel® Cluster Ready version compliance	Enabled
imb_pinpgong	MPI performance	Enabled
infiniband	InfiniBand* uniformity and wellness	Enabled
intel_cluster_runtime	Intel® Cluster Ready runtime library compliance	Enabled
iozone	Disk I/O performance	Enabled
java	Java* uniformity and functionality	Enabled
kernel	Linux* kernel	Enabled
kernel_param	Kernel parameter uniformity	Enabled
libraries	Intel® Cluster Ready runtime library compliance	Enabled
lshw	Hardware uniformity	Enabled
memory	Memory compliance	Enabled
miccheck	Intel® Xeon Phi™ coprocessor uniformity and wellness	Enabled
micinfo	Intel® Xeon Phi™ coprocessor uniformity and wellness	Enabled
mount	Mount point compliance and uniformity	Enabled
mpi_internode	Multi-node MPI functionality	Enabled
mpi_local	Single-node MPI functionality	Enabled
ntp	Clock synchronization	Enabled
offload_phi	Intel® Xeon Phi™ coprocessor functionality	Enabled
opa	Intel® OPA uniformity and wellness	Enabled

Check	Description	Default State
perl	Perl* compliance, uniformity, and functionality	Enabled
process	Process table	Enabled
python	Python* compliance, uniformity, and functionality	Enabled
rpm	RPM uniformity	Enabled
shells	Shell compliance	Enabled
storage	Disk capacity	Enabled
stream	Memory bandwidth performance	Enabled
tcl	Tcl compliance, uniformity, and functionality	Enabled
x11_tools	X11 tool compliance	Enabled